Hiten Shah: when AI builds anything, the PM's job is deciding what deserves the work

PM Daily Digest

2026-11-08

Hiten Shah: when AI builds anything, the PM's job is deciding what deserves the work

By PM Daily Digest • November 8, 2026

Hiten Shah argues that cheap AI building makes testing assumptions the real measure of productivity. Also: Replit's customer-observation origin story, a PM community split over whether PMs must own revenue, and skepticism about AI-run win/loss interviews.

Finished work isn't validated work

In a new essay, Hiten Shah recalls the "Hiten Bombs" he used to drop at KISSmetrics: customer comments or competitor features he shared that made people stop their work to check whether priorities had changed. Eventually the head of product wrote a 1,300-word memo about the disruption, and the team was "losing the focus that had given us our lead" [1]. Today a coding agent can turn the same kind of idea into working software "before I've decided whether it's worth pursuing" [1].

His concern is a flood of competent work built on assumptions nobody has questioned [1]. He cites a 2026 study of cheap consumer-product ideas. AI-generated ideas got higher predicted purchase-intent ratings and were seven times more likely to land in the top tenth, but the AI pool was less diverse. These were evaluations of ideas, not market tests [1]. Generating larger pools and combining models helped widen the range [1].

KISSmetrics' own history shows what goes wrong. Its first users, Facebook app developers, liked the product but couldn't pay. Its next product, a dashboard, asked customers to do too much setup. Momentum came only after customer research turned up online marketers who would pay for answers to specific questions [1]. His proposed metric: judge productivity by the uncertainty you remove. After each experiment, ask which assumption you tested

and what changed about your next decision [1]. In a separate post he wrote that "curiosity doesn't like per-token pricing" [2].

Replit: watch users, then automate the workaround

On YC's channel, Amjad Masad told Paul Graham that Replit's agent idea came from inviting users to the office and watching them. These were people who had never coded before. They pasted ChatGPT code into Replit, copied the errors back into ChatGPT, and pasted the fixes back in [3]. Automating that loop meant understanding the errors and repeating until the program worked, rather than stopping after one turn [3].

Other points from the interview: - **Strategy.** Resist announcing a big strategy. "Let the customers tell the story": help one customer do something, then make it available to everyone [3]. - **Enterprise.** Security guardrails sometimes make Replit slower and more expensive, but enterprises value them [3]. - **AI budgets.** Masad disagrees with giving every employee the same monthly token allowance. Buyers should look at ROI, because one person can produce 100x what another does. In the short term, vendors may have to provide the controls buyers ask for, then build better ways to show where the value is [3].

Should every PM own revenue?

A high-traffic r/ProductManagement post argued that many PMs can't explain why customers buy, why deals die or where revenue leaks, and are weak on unit economics, pricing and segmentation. "Ignorance of revenue is neither a product nor a career strategy" [4].

For: One commenter prioritizes by estimated revenue impact and says even bugs, compliance and infrastructure margins can be tied to revenue [5]. An interim PM who interviewed 250+ companies said fewer than 6 had anything like a product strategy or value model. In 9 of 10 projects, senior management showed no interest in revenue and rewarded predictable output [6].

Against: In big tech, infrastructure, privacy or engagement-focused PMs may never touch revenue directly [7]. In deep enterprise org charts, separate commercial or strategy teams own the business case, and line PMs mainly execute [8]. One founder built a tool to measure the revenue impact of customer pain points. After 60+ interviews found PMs weren't interested, the company pivoted to selling to Customer Success [9].

Market research tactics

- **AI win/loss interviews** (Clozd, Klue): a PMM worried customers won't talk to an AI and that it could hurt the company's reputation [10]. Another said a previous employer's rollout "didn't go over well" [11]. One commenter instead runs simulated interviews, five buying-committee per-

sonas across 100 simulated companies, to prepare for objections at customer advisory board meetings [12].

- **Turning a custom build into a product:** a PMM who built an AI sales-enablement hub for one company asked whether to sell it to others [13]. Advice:
  - Interview 5–10 similar teams to see whether they share inputs, approvals and CRM structure.
  - Sell a narrow paid pilot, and track setup hours, exceptions and usage.
  - If every deployment needs heavy custom mapping, start as a productized service [14].
  - Competitors include Highspot, Seismic, Mindtickle and Showpad, so the suggested angle is an affordable option for SMBs [15].
- **PM hiring:** a r/ProductMgmt thread asked whether interviews should weigh how candidates reason through ambiguous situations more than how fluently they apply frameworks [16].

Sources

1. X article by @hnshah
2. X post by @hnshah
3. How To Build The Future: Paul Graham & Amjad Masad
4. r/ProductManagement post by u/DeanOnDelivery
5. r/ProductManagement comment by u/Exotic-Sale-3003
6. r/ProductManagement comment by u/Bernhard-Welzel
7. r/ProductManagement comment by u/pizza_the_mutt
8. r/ProductManagement comment by u/Reginherus
9. r/ProductManagement comment by u/wifflebal
10. r/ProductMarketing post by u/SquidwardDab
11. r/ProductMarketing comment by u/Used-Media5980
12. r/ProductMarketing comment by u/monkeythumpa
13. r/ProductMarketing post by u/Comfortable_Tax2029
14. r/ProductMarketing comment by u/macromind
15. r/ProductMarketing comment by u/BangPC
16. r/ProductMgmt post by u/Aware_Supermarket_34